Mutual Information and Conditional Mean Prediction Error

Clive G. Bowsher[1] and Margaritis Voliotis
School of Mathematics,
University of Bristol, U.K.

arXiv:1407.7165v1 [cs.IT] 26 Jul 2014

Abstract

Mutual information is fundamentally important for measuring statistical dependence between variables and for quantifying information transfer by signaling and communication mechanisms. It can, however, be challenging to evaluate for physical models of such mechanisms and to estimate reliably from data. Furthermore, its relationship to better known statistical procedures is still poorly understood. Here we explore new connections between mutual information and regression-based dependence measures, ν^{-1}, that utilise the determinant of the second-moment matrix of the conditional mean prediction error. We examine convergence properties as $\nu \to 0$ and establish sharp lower bounds on mutual information and capacity of the form $\log(\nu^{-1/2})$. The bounds are tighter than lower bounds based on the Pearson correlation and ones derived using average mean square-error rate distortion arguments. Furthermore, their estimation is feasible using techniques from nonparametric regression. As an illustration we provide bootstrap confidence intervals for the lower bounds which, through use of a composite estimator, substantially improve upon inference about mutual information based on k-nearest neighbour estimators alone.

Index Terms

Lower bound on mutual information, relative entropy, information capacity, nearest-neighbour estimator, correlation and dependence measures, regression.

I. Introduction

Mutual information is fundamentally important for measuring statistical dependence between variables [1], [2], [3], [4], and for quantifying information transfer by engineered or naturally occurring communication systems [5], [6]. Statistical analysis using mutual information has been particularly influential in neuroscience [7], and is becoming so in systems biology for studying the biomolecular signaling networks used by cells to detect, process and act upon the chemical signals they receive [8], [9], [10]. It can, however, be challenging to estimate mutual information reliably with available sample sizes [11], and difficult to derive mutual information and capacity exactly using mechanistic models of the 'channels' via which signals are conveyed. Furthermore, connections between mutual information and better known statistical procedures such as regression, and their associated dependence measures, are still poorly understood. In order to address these challenges, the relationship between mutual information and the error incurred by estimation (or 'prediction') using the conditional mean is now receiving attention. The focus has been on minimum mean square estimation error or, more generally, its average across the elements of the vector being estimated [12], [13], [14]. Instead, we focus on connections between mutual information and regression-based dependence measures, ν^{-1}, that utilise the determinant of the second-moment matrix of the conditional mean prediction error. We examine convergence properties as $\nu \to 0$, and establish sharp lower bounds on mutual information of the form $\log(\nu^{-1/2})$. The bounds are tighter than lower bounds based on the Pearson correlation and ones derived using average mean square-error rate distortion arguments.

The *mutual information* between 2 random vectors X and Z, written $I(X;Z)$, is the Kullback-Leibler divergence between their joint distribution and the product of their marginal distributions [15]. Mutual information thus measures the divergence between the joint distribution of (X, Z) and the distribution in which X and Z are independent but have the same marginals. $I(X;Z)$ has desirable properties as a measure of statistical dependence: it satisfies (after monotonic transformation) all 7 of Rényi's postulates [16] for a dependence measure between 2 random variables, and underlies the recently introduced maximal information coefficient [4] for detecting pairwise associations in large data sets. Importantly, mutual information captures nonlinear dependence and dependence arising from moments higher than the conditional mean.

A decision-theoretic interpretation is indicative of the broad applicability of mutual information as a summary measure of statistical dependence. It can be shown [17] that $I(X;Z)$ is equal to the increase in expected utility from

[1] corresponding author: c.bowsher@bristol.ac.uk

reporting the posterior distribution of either of the two random vectors based on observation of the other, compared to reporting its marginal distribution—for example, reporting $p(Z|X = x)$ instead of $p(Z)$. This holds when the utility function is a smooth, proper, local score function—as appropriate for scientific reporting of distributions as 'pure inferences' [18]—because the logarithmic score function is the only score function having all these properties. In information theory, the supremum of $I(X;Z)$ over the set of allowed input distributions $\mathcal{F}$, termed the *information capacity*, equates to the maximal errorless rate of information transmission over a noisy channel when the channel is used for long times [5].

The above discussion makes clear that, from a rich variety of perspectives, mutual information is fundamentally important for measuring statistical dependence and for quantifying information transfer by signaling and communication mechanisms.

Our general setting may be depicted as

$$X \to Y \to Z, \tag{1}$$

with (X, Y, Z) a real-valued random vector. Here Z is conditionally independent of X given Y. The conditional distribution of Z given X is determined by some physical mechanism, whose 'internal' variables are denoted by Y in Eq. 1. Such a mechanism is often termed a channel in information theory, although we do not restrict attention to signaling and communication channels here. Often we have in mind situations where X causes Z but not *vice versa*, and the conditional distribution of Z given X does not depend on the experimental 'regime' giving rise to the distribution of X [19]. There is always an asymmetry between X and Z in the general setting we consider. We term X the *input* or *treatment* because its marginal distribution can in principle be any distribution (although we may wish to restrict attention to particular classes thereof). In contrast, the *output* or *response* Z is the realisation of the mechanism given the input X. In general, not all marginal distributions for Z can be obtained for a given mechanism by appropriate choice of the marginal of X. When analysing the probabilistic properties of physical models of mechanisms, the distribution (or set of distributions) for the input X is given, but the marginal distribution of Z is often unknown. In experimental settings, the input distribution is taken not to affect the conditional distribution of Z given X, or else can sometimes be directly manipulated.

Examples of our general setting include experimental design with X as the treatment and Z the response of interest; and signaling or communication channels with X as the input signal and Z its noisy representation. An example of a scientific area of application is the current effort to understand the biomolecular signaling mechanisms used by living cells to relay the chemical signals they receive from their environment [8]. Here, the interest is both in understanding why some biomolecular mechanisms perform better than others, and in measuring experimentally in the laboratory the mutual information between X and Z or the information capacity. Broadly speaking, the first involves deriving dependence measures between X and Z for different stochastic mechanisms (given a particular input distribution). The second might involve, for example, nonparametric estimation of the mutual information between the concentration of a chemical treatment applied to the cells and the level of an intracellular, biochemical output.

We note, however, that the formal statements of our results do not require any particular interpretation of X and Z. Rather, the general setting just described motivates the results and places them in context. A sequential reading of Equations 2 to 12 inclusive provides a convenient preview of our theoretical results establishing lower bounds on mutual information and information capacity.

II. Setup and Notation

For random vectors X and Z, we define $\nu(Z|X) = \det\left(\mathbb{E}\{\mathbb{V}\left[Z|X\right]\}\right)/\det\left(\mathbb{V}\left[Z\right]\right)$, where $\mathbb{V}$ denotes a covariance matrix. In general, $\nu(Z|X)$ is not equal to $\nu(X|Z)$. For 2 scalar random variables, $\nu(Z|X)$ is equal to the minimum mean square estimation error or minimum MSEE for estimation of Z using X, normalised by the variance of Z (since $\mathbb{E}\left[Z|X\right]$ is the optimal estimator). We denote the optimal estimation or 'prediction' error by $e(Z|X) = Z - \mathbb{E}\left[Z|X\right]$. In general, $\nu(Z|X)$ is the ratio of the determinant of the second-moment matrix of the error $e(Z|X)$ and the determinant of the variance matrix of Z, that is

$$\nu(Z|X) = \frac{\det\left(\mathbb{E}\left[e(Z|X)e(Z|X)^{\mathrm{T}}\right]\right)}{\det\left(\mathbb{V}\left[Z\right]\right)}. \tag{2}$$

We will show that $\nu(Z|X)^{-1}$ provides a generalised measure of 'signal-to-noise', applicable to non-Gaussian settings, that relies on first and second conditional moments of Z given X (via the law of total variance) rather than on all features of the joint distribution. We make few assumptions about the conditional density describing the mechanism or channel $f(z|x)$, except that the conditional mean $m(x) = \mathbb{E}[Z|X = x]$ is an invertible, continuously differentiable function of x. A central result of the paper (see Theorem 5 and Corollary 6) is then that

$$I(X;Z) \geq \log\left\{\nu(Z|X)^{-1/2}\right\} \geq \frac{-d_Z}{2}\log\left\{\frac{d_Z^{-1}\mathrm{tr}(\mathbb{E}\{\mathbb{V}[Z|X]\})}{[\det(\mathbb{V}[Z])]^{d_Z^{-1}}}\right\}, \tag{3}$$

where all terms are evaluated under the joint density for (X, Z) implied by the channel $f(z|x)$ and a Gaussian density for the transformed input, $m(X)$. Here d_Z is the dimension of the vector Z. The second term in Eq. 3 is our lower bound utilising the determinant of the second-moment matrix of the prediction error of the conditional mean $\mathbb{E}[Z|X]$, while the third term instead utilises the average mean square error of that conditional mean. We discuss the relation of the third term to rate distortion arguments later in the paper. Notice that characterising the first and second conditional moments of the mechanism, $\mathbb{E}[Z|X]$ and $\mathbb{V}[Z|X]$, is enough (via the law of total variance) to evaluate the lower bound $\log\left\{\nu(Z|X)^{-1/2}\right\}$ for a given Gaussian distribution of $m(X)$. Maximising the bound over such distributions then also yields a useful bound on the information capacity.

As a first step in analysing connections between mutual information and our regression-based measures, we explore the relationship between the convergence to zero of $\nu(Z|X)$ or $\nu(X|Z)$, and the convergence of mutual information. For simplicity, we analyse the bivariate case where the variable X has finite support, for example a finite collection of treatment concentrations in a cell signaling experiment. We write I for mutual information, H for discrete entropy and h for differential entropy.

III. Convergence properties

Theorem 1. *Let (X_n, Z_n) be a sequence of pairs of real-valued random variables, with the support of X_n given by a finite set $\mathcal{X}_n$ ($|\mathcal{X}_n| \geq 2$ and bounded above by a constant $\forall n$). Write $m_n(X_n)$ for the function $\mathbb{E}[\breve{Z}_n|X_n]$, where $\breve{Z}_n \triangleq Z_n\mathbb{V}[Z_n]^{-\frac{1}{2}}$. Denote its support by $\mathcal{M}_n = \{m_n(x); x \in \mathcal{X}_n\} \subset \mathbb{R}$. Let ϵ_n^* be 1/2 of the minimum distance between any two points in $\mathcal{M}_n$ and $\epsilon^* \triangleq \inf\{\epsilon_n^*\}$. Suppose that:*

(1) $\nu(Z_n|X_n) \to 0$ as $n \to \infty$; (2) the functions m_n are one-to-one mappings from $\mathcal{X}_n$ to the real line; and (3) $\epsilon^* > 0$, so that any pair of points in a support $\mathcal{M}_n$ are separated by at least $2\epsilon^*$. Then as $n \to \infty$, $H(X_n|Z_n) \to 0$ and therefore $H(X_n) - I(X_n; Z_n) \to 0$.

In Theorem 1, the response variable Z is real-valued and can, for example, be either a continuous or discrete random variable. Theorem 1 establishes the convergence of the mutual information $I(X;Z)$ to the entropy of X as $\nu(Z|X) \to 0$, under the condition that the conditional mean $\mathbb{E}[Z|X]$ is an invertible function of X. By definition, $I(X;Z)$ cannot be greater than the entropy of X. The intuition for the result in Theorem 1 is that the convergence of $\nu(Z|X)$ to zero enables construction of a point estimator of X (utilising the conditional expectation $\mathbb{E}[Z|X]$) whose performance becomes 'perfect' in this limit. The condition requiring invertibility of the conditional mean function would be needed even in the case where Z is a deterministic function of X, otherwise $I(X;Z) \leq H(Z) < H(X)$.

We have rescaled the output Z_n so that $\mathbb{V}[\breve{Z}_n]$ is constant at 1 for all n. In particular, Theorem 1 does *not* require the minimum mean square estimation error, $\mathbb{E}\{(Z - \mathbb{E}[Z|X])^2\}$, to converge to zero as $n \to \infty$. A physical example where the minimum MSEE diverges but Theorem 1 applies is given by a molecular signaling system, with Z the number of output molecules, which is operating in the macroscopic (or large system size) limit where the dynamics of chemical concentrations are deterministic, conditional on the input X.

Proof: We have that $\nu(Z_n|X_n) = \nu(\breve{Z}_n|X_n) \to 0$. Since $\nu(\breve{Z}_n|X_n) = \mathbb{E}\{\mathbb{V}[\breve{Z}_n|X_n]\} = \mathbb{E}\{(\breve{Z}_n - \mathbb{E}[\breve{Z}_n|X_n])^2\}$, it follows that $\breve{Z}_n - \mathbb{E}[\breve{Z}_n|X_n]$ converges to zero in mean square (in L^2) and therefore $\breve{Z}_n - \mathbb{E}[\breve{Z}_n|X_n] \to^{pr} 0$ (where $\to^{pr}$ denotes convergence in probability). Since $H(X_n|\breve{Z}_n) = H(X_n|Z_n) = H(X_n) - I(X_n, Z_n)$, we must establish that $H(X_n|\breve{Z}_n) \to 0$ as $n \to \infty$. Consider estimating X_n based on observation of $\breve{Z}_n$ using the estimator, $\widehat{X}_n(\breve{Z}_n)$, which is equal to the a point $x \in \mathcal{X}_n$ that minimises the distance $|\breve{Z}_n - m_n(x)|$. Condition (3) applies not to $\mathbb{E}[Z_n|X_n = x]$ but to $\mathbb{E}[\breve{Z}_n|X_n = x]$. Let $z \in \mathbb{R}, m_n \in \mathcal{M}_n$ and notice that if $|z - m_n| < \epsilon^*$, then $|z - m_n| < \epsilon_n^* < |z - m_n'|$, that is m_n is closer to z than is any other point m_n' in $\mathcal{M}_n$. Therefore, if $|\breve{Z}_n - \mathbb{E}[\breve{Z}_n|X_n]| < \epsilon^*$, then

$\mathbb{E}[\breve{Z}_n|X_n]$ is located at the unique point $m_n \in \mathcal{M}_n$ that minimises $|\breve{Z}_n - m_n|$ and, using condition (2), $\widehat{X}_n(\breve{Z}_n) = X_n$, that is, the estimator recovers X_n without error. The probability of estimation error, p_{error}, satisfies

$$p_{\text{error}} = \mathbb{P}\left\{\hat{X}_n\left(\breve{Z}_n\right) \neq X_n\right\} \leq \mathbb{P}\left\{|\breve{Z}_n - \mathbb{E}[\breve{Z}_n|X_n]| \geq \epsilon^*\right\},$$

and therefore $p_{\text{error}} \to 0$ as $n \to \infty$. Fano's Inequality [15] then implies that $H(X_n|\breve{Z}_n) \to 0$ as required. ∎

We have thus shown in the context of Theorem 1 that the limit $\nu(Z|X) \to 0$ characterises a regime of large signal-to-noise for the input X, without the need to impose further conditions on the joint distribution of X and Z. In [12], the authors consider the opposite regime of low signal-to-noise for non-linear channels with additive Gaussian noise. They obtain an asymptotic expansion of the mutual information, $I(X;Z)$, whose leading term is a decreasing function of a variable which, in their setting, is equal to $\nu(Z|X)$. In Theorem 2 below, we consider the case where the conditioning is on the response Z instead of on the input X (see Section V for further discussion of regression on the response variable).

Theorem 2. *Let (X_n, Z_n) be a sequence of pairs of real-valued random variables, with the support of X_n given by a finite set $\mathcal{X}_n$ ($|\mathcal{X}_n| \geq 2$ and bounded above by a constant $\forall n$). Define $\breve{X}_n \triangleq X_n \mathbb{V}[X_n]^{-\frac{1}{2}}$, with support $\breve{\mathcal{X}}_n$, and let ϵ_n^* be 1/2 of the minimum distance between any two points in $\breve{\mathcal{X}}_n$. Suppose that $\epsilon^* \triangleq \inf\{\epsilon_n^*\} > 0$. If $\nu(X_n|Z_n) \to 0$ as $n \to \infty$, then $H(X_n|Z_n) \to 0$ and $H(X_n) - I(X_n;Z_n) \to 0$.*

Proof: Given in the Appendix, using an argument similar to the proof of Theorem 1. ∎

Again, the mutual information $I(X;Z)$ converges to the entropy of X, this time as $\nu(X|Z) \to 0$. (We note in passing that, if both X and Z have finite support, a corollary of Theorem 2 is that: either $H(X_n) - H(Z_n) \to 0$, or $\nu(X_n|Z_n)$ and $\nu(Z_n|X_n)$ do not simultaneously converge to zero).

Theorems 1 and 2 establish connections between regression-based measures of dependence ν^{-1} and mutual information, without imposing strong assumptions about the joint distribution of (X, Z). We conjecture that similar theorems will hold for random variables X and Z having a joint density with respect to Lebesgue measure. These theorems indicate the possibility, explored below, of lower bounding the mutual information using ν^{-1}. A consequence of our Theorem 5 is that, for the general multivariate, absolutely continuous case, $I(X;Z) \to \infty$ when $\nu(Z|X) \to 0$ (and $\mathbb{E}[Z|X = x]$ is an invertible mapping), and ν is evaluated under the appropriate marginal distribution for the input X.

IV. Lower bounds on mutual information and capacity

Our aim in this and the subsequent section is to establish lower bounds on mutual information, $I(X;Z)$, for an absolutely continuous random vector (X, Z) with finite dimension $d \geq 2$. These bounds will hold under certain marginal distributions for the input X, and also provide lower bounds on the information capacity. We first show that the following Lemma holds when the marginal distribution of X is normal. The Lemma relates the mutual information of X and Z to their variances and covariance.

Lemma 3. *Let (X, Z) be a random vector in $\mathbb{R}^d$, $d \geq 2$, with joint density $f(x, z)$ with respect to Lebesgue measure. Suppose the density of X, $f(x)$, is multivariate normal and that (X, Z) has finite variance matrix under f. Then*

$$\begin{aligned} I_f(X;Z) &\geq \log\left\{\frac{\det(\Sigma_{XX})}{\det(\Sigma_{XX} - \Sigma_{XZ}\Sigma_{ZZ}^{-1}\Sigma_{ZX})}\right\}_f^{1/2} \\ &= \log\left\{\frac{\det(\Sigma_{ZZ})}{\det(\Sigma_{ZZ} - \Sigma_{ZX}\Sigma_{XX}^{-1}\Sigma_{XZ})}\right\}_f^{1/2} \end{aligned} \tag{4}$$

where, for example, $\Sigma_{ZX} = \text{Cov}(X, Z) = \mathbb{E}\{(X - \mathbb{E}[X])(Z - \mathbb{E}[Z])^{\text{T}}\}$, and subscript f indicates that the mutual information and the covariance matrices are those under the joint density $f(x, z)$.

Proof: Given in the Appendix. ∎

When $d = 2$ (and, again, $f(x)$ is normal), Lemma 3 simplifies to give the lower bound

$$I(X;Z) \geq \log\{[1 - \text{Corr}^2(X, Z)]^{-1/2}\},$$

where Corr denotes the Pearson correlation. This inequality shows how to relate perhaps the best known measure of association to mutual information. An outline proof is given for the $d = 2$ case in [20]. The multivariate, $d > 2$, case has not to our knowledge appeared previously, although it is a reasonably straightforward generalisation. We provide a complete proof for $d \geq 2$ in the Appendix.

Remark 4. The lower bounds in Eq. 4 are the mutual information of a multivariate normal density, g, with identical covariance matrix to that of f, the true joint density of X and Z. We want to obtain a lower bound for $I_f(X;Z)$ in terms of $\nu_f(Z|X)$. It might appear that one way to do so would be to adopt a similar strategy, attempting to show that $I_f(X;Z)$ is bounded below by the mutual information of a multivariate normal, g, having identical $\nu(Z|X)$ to f (and with the marginal density $g(x) = f(x)$). However, this strategy fails. Although $I_g(X;Z)$ still depends only on $\nu(Z|X)$, the equality $\mathbb{E}_f\{\log[g(X,Z)/f(X)g(Z)]\} = \mathbb{E}_g\{\log[g(X,Z)/f(X)g(Z)]\}$ no longer holds in general (see Eq. 13 and subsequent argument in the Appendix).

Our general strategy to obtain lower bounds on mutual information and capacity is as follows. First, we specify a channel with suitable 'pseudo-output' and Gaussian 'pseudo-input'. These are transformed versions of Z and X respectively (sometimes we transform X alone). The transformations are given by the conditional mean functions: for example, the pseudo-input may be $m(X) = \mathbb{E}[Z|X]$, as in Theorem 5 below. Second, we apply Lemma 3 with the pseudo-input and pseudo-output in place of X and Z there. We then make use of the following relationship that holds for any random vector (U, W):

$$\begin{aligned} \mathrm{Cov}(U, \mathbb{E}[U|W]) &= \mathbb{E}\left\{\mathbb{E}\left[(U - \mathbb{E}[U])(\mathbb{E}[U|W] - \mathbb{E}[U])^{\mathrm{T}}|W\right]\right\} \\ &= \mathbb{V}\{\mathbb{E}[U|W]\}. \end{aligned} \tag{5}$$

We now use this strategy to obtain a lower bound for the mutual information, $I_f(X;Z)$, in terms of $\nu_f(Z|X)$.

Theorem 5. *Let (X, Z) be a random vector in $\mathbb{R}^d, d \geq 2$. Consider the conditional density (with respsect to Lebesgue measure), $f(z|x)$, and suppose that $m(x) = \mathbb{E}[Z|X = x]$ is a one-to-one, continuously differentiable mapping (whose domain is an open set and a support of X). Let $m(X)$ be normally distributed and denote by $f(x)$ the implied density of X. Then*

$$I_f(X;Z) \geq \log\left\{\nu_f(Z|X)^{-1/2}\right\}, \tag{6}$$

where we assume moments are finite such that $\nu_f(Z|X)$ is well defined and that $\mathbb{E}\{\mathbb{V}[Z|X]\}$ is non-singular under the joint density, f. This lower bound is sharp since $I_f(X;Z) = \log\left\{\nu_f(Z|X)^{-1/2}\right\}$ when $f(x,z)$ is a multivariate normal density. Furthermore, the information capacity $C_{Z|X}$ satisfies

$$C_{Z|X} \geq \log\left\{\nu_f(Z|X)^{-1/2}\right\}, \tag{7}$$

provided $f(x)$ is the density of an allowed input distribution, that is a distribution in $\mathcal{F}$.

Proof: Consider the mechanism $M \to X \to Z$, where $X = m^{-1}(M)$ and X is then applied to the channel $f(Z|X)$. Here $M = m(X)$ is the transformed or pseduo-input. Notice that $\mathbb{E}[Z|M] = \mathbb{E}[Z|X] = M$ because $\sigma(X) = \sigma(M)$, and therefore $\mathrm{Cov}(M,Z) = \mathrm{Cov}(Z, \mathbb{E}[Z|M]) = \mathbb{V}\{\mathbb{E}[Z|M]\} = \mathbb{V}[M]$, by Eq. 5 (and using twice that $\mathbb{E}[Z|M] = M$). Under the conditions of Theorem 5, M has a Gaussian distribution, with the distribution of X that implied by the mapping $X = m^{-1}(M)$. Given the properties of the one-to-one mapping $m(x)$, we also have that $I_f(M;Z) = I_f(X;Z)$ (see, for example, [21]). Applying Lemma 3 to the Gaussian, pseudo-input M and the output Z yields

$$\begin{aligned} I_f(X;Z) &\geq \frac{1}{2}\log\left\{\frac{\det(\mathbb{V}[Z])}{\det(\mathbb{V}[Z] - \mathrm{Cov}(Z,M)\mathbb{V}[M]^{-1}\mathrm{Cov}(Z,M)^{\mathrm{T}})}\right\}_f \\ &= \frac{1}{2}\log\left\{\frac{\det(\mathbb{V}[Z])}{\det(\mathbb{V}[Z] - \mathbb{V}\{\mathbb{E}[Z|M]\})}\right\}_f, \end{aligned}$$

where the second line again uses $\mathrm{Cov}(Z,M) = \mathbb{V}[M] = \mathbb{V}\{\mathbb{E}[Z|M]\}$. The result then follows directly because $\mathbb{V}[Z] = \mathbb{V}\{\mathbb{E}[Z|M]\} + \mathbb{E}\{\mathbb{V}[Z|M]\}$, and because we have that $\mathbb{V}[Z|M] = \mathbb{V}[Z|X]$ since $\sigma(M) = \sigma(X)$. Eq. 7 follows from Eq. 6 and the definition of capacity as the supremum of mutual information over the collection of allowed input distributions, $\mathcal{F}$. ■

Notice that characterising the first and second conditional moments of the mechanism, $\mathbb{E}[Z|X]$ and $\mathbb{V}[Z|X]$, is enough (using the law of total variance) to evaluate the lower bound $\log\left\{\nu(Z|X)^{-1/2}\right\}$ for a given Gaussian distribution of $m(X)$. Maximising the bound over such distributions then yields the largest lower bound on the information capacity. The approach is applicable when experimental data have been generated under some other input distribution, provided the first and second conditional moments are carefully estimated.

We now discuss the relationship of our lower bounds in Eqs. 6 and 7 to average mean-square error, rate distortion-type arguments with a Gaussian source.

Corollary 6. *Let d_Z be the dimension of Z and suppose that the conditions of Theorem 5 apply. It follows from Eq. 6, scaling by the reciprocal of d_Z, that*

$$\begin{aligned} d_Z^{-1} I_f(X;Z) \geq \frac{1}{2} d_Z^{-1} \log\{\nu_f(Z|X)^{-1/2}\} &= \frac{1}{2}\log\left\{\frac{[\det(\mathbb{V}[Z])]^{d_Z^{-1}}}{[\det(\mathbb{E}\{\mathbb{V}[Z|X]\})]^{d_Z^{-1}}}\right\}_f \\ &\geq \frac{1}{2}\log\left\{\frac{[\det(\mathbb{V}[Z])]^{d_Z^{-1}}}{d_Z^{-1}\mathrm{tr}(\mathbb{E}\{\mathbb{V}[Z|X]\})}\right\}_f, \end{aligned} \tag{8}$$

since $\mathbb{E}_f\{\mathbb{V}[Z|X]\}$ is positive definite, where f indicates evaluation under the joint density $f(x,z)$. We have used that $0 < [\det(M)]^{1/m} \leq m^{-1}\mathrm{tr}(M)$ for a positive definite, $m \times m$ matrix, M [15].

Notice that $d_Z^{-1}\mathrm{tr}(\mathbb{E}\{\mathbb{V}[Z|X]\})$ is the average minimum MSEE given X, the average being across the scalar components of the vector Z. Eq. 8 establishes that our lower bound, $\log\{\nu_f(Z|X)^{-1/2}\}$, is tighter than the lower bound based on the average minimum MSEE, $\frac{1}{2}\log\{\det(\mathbb{V}[\mathrm{Z}])^{d_Z^{-1}} / \; d_Z^{-1}\mathrm{tr}(\mathbb{E}\{\mathbb{V}[Z|X]\})\}_f$. The two lower bounds clearly coincide for the bivariate case, $d = 2$. The bound based on the average minimum MSEE might, at first sight, appear to have the form that would be obtained by a rate distortion-type argument [22, Section 4.5.2] with Z as the Gaussian 'source'. However, this is not the case because the bound would in general be evaluated under a joint density for (X, Z) different than f. (However, see Lemma 8 and the subsequent discussion for the case of regression on the response variable). Notice also that in our setting of Eq. 1, we may not be able to adjust the input distribution in order to obtain a Gaussian marginal for Z.

Instead, one might treat $M = \mathbb{E}[Z|X]$ as the Gaussian source in a rate distortion-type argument. Consider the case $d = 2$. One obtains the result $I_f(X;Z) \geq \frac{1}{2}\log(\mathbb{V}\{\mathbb{E}[Z|X]\}/\mathbb{E}\{\mathbb{V}[Z|X]\})_f$, where the numerator is the variance of the source and the denominator is the expected square-error distortion between the source and its estimate, here Z. The right-hand side of this inequality is strictly less than our lower bound, $\log\{\nu_f(Z|X)^{-1/2}\} = \frac{1}{2}\log(\mathbb{V}[Z]/\mathbb{E}\{\mathbb{V}[Z|X]\})_f$, since $\mathbb{V}[Z] - \mathbb{V}\{\mathbb{E}[Z|X]\} = \mathbb{E}\{\mathbb{V}[Z|X]\} > 0$. An analogous argument applies to the case $d > 2$, since $\det(\mathbb{V}[Z]) > \det(\mathbb{V}\{\mathbb{E}[Z|X]\})$. (For the case where X itself is Gaussian, see Lemma 8). We conclude that our lower bounds in Eqs. 6 and 7 are tighter than lower bounds derived using average mean-square error rate distortion-type arguments with a Gaussian source.

V. Regression on the response variable

We have so far considered lower bounds on mutual information that rely on the error in estimation of the response variable Z, using the conditional mean of Z given X. In this section we instead consider lower bounds that utilise the regression of X on the response variable, Z. Using the novel proof strategy adopted for Theorem 5, we can show that the lower bound $\log\{\nu_f(\tilde{X}|Z)^{-1/2}\}$ improves upon the one based on the Pearson correlation, $\log\{[1 - \mathrm{Corr}_f^2(\tilde{X}, Z)]^{-1/2}\}$, where $\tilde{X}$ is the result of transforming the input to have a Gaussian marginal distribution. We can thus always (weakly) improve upon the lower bound for the bivariate case given by Lemma 3. Intuitively, the improvement arises because $\nu_f(\tilde{X}|Z)^{-1}$ captures dependence from non-linearity in the conditional mean, $\mathbb{E}[\tilde{X}|Z]$, whereas the Pearson correlation does not.

Theorem 7. *Let (X, Z) be a random vector in $\mathbb{R}^2$ with joint density $f(x,z)$ with respect to Lebesgue measure. Suppose there exists a one-to-one mapping $s : X \to \tilde{X}$ (whose domain is an open set and a support of X) such that $\tilde{X}$ has a Gaussian density. We assume that the mapping s is continuously differentiable with derivative that is everywhere non-zero, and that $\nu_f(\tilde{X}|Z)^{-1/2}$ exists. Then*

$$I_f(X;Z) \geq \log\{\nu_f(\tilde{X}|Z)^{-1/2}\} \geq \log\{[1 - \mathrm{Corr}_f^2(\tilde{X}, Z)]^{-1/2}\}, \tag{9}$$

where the third term is the lower bound on $I_f(\tilde{X};Z) = I_f(X;Z)$ given by Lemma 3 in the case $d = 2$. We assume $\mathrm{Corr}_f(\tilde{X}, Z)$ is well defined and less than 1. Subscript f indicates evaluation under the joint distribution implied by $f(x, z)$.

As we show in the proof below, the lower bound $\log\{\nu_f(\tilde{X}|Z)^{-1/2}\}$ can be understood as the result of first transforming Z to the 'pseudo-output' $t(Z) = \mathbb{E}[\tilde{X}|Z]$, and then basing the bound on $\mathrm{Corr}_f^2(\tilde{X}, t(Z))$, using Lemma 3. We then establish that using another (measurable) transformation, $t(Z)$, cannot yield a greater bound (given some choice of the Gaussian variable $\tilde{X}$). This includes, in particular, the lower bound based on the squared correlation of $\tilde{X}$ and Z itself. Notice that we cannot construct a lower bound using the maximal correlation of (X, Z) [23], [16], because the implied transformation of X need not result in the Gaussian distribution needed to apply Lemma 3.

Proof: Given the properties of the one-to-one mapping $s(x)$, we have that $I_f(X;Z) = I_f(\tilde{X};Z)$. Consider the mechanism $\tilde{X} \to X \to Z \to \mathbb{E}[\tilde{X}|Z]$, in which we first transform $\tilde{X}$ to X and then apply X to the 'channel' $f(Z|X)$. The pseudo-output here is $\mathbb{E}[\tilde{X}|Z]$. By the data processing inequality, $I_f(\tilde{X};Z) \geq I_f(\tilde{X};\mathbb{E}[\tilde{X}|Z])$. Applying Lemma 3 to the Gaussian pseudo-input $\tilde{X}$ and the pseudo-output $\mathbb{E}[\tilde{X}|Z]$ yields

$$I_f(\tilde{X};\mathbb{E}[\tilde{X}|Z]) \geq \frac{1}{2}\log\{[1 - \mathrm{Corr}^2(\tilde{\mathrm{X}}, \mathbb{E}[\tilde{\mathrm{X}}|\mathrm{Z}])]^{-1/2}\}.$$

By Eq. 5, $\mathrm{Cov}(\tilde{X}, \mathbb{E}[\tilde{X}|Z]) = \mathbb{V}\{\mathbb{E}[\tilde{X}|Z]\}$. Hence $\mathrm{Corr}^2(\tilde{X}, \mathbb{E}[\tilde{X}|Z]) = \mathbb{V}\{\mathbb{E}[\tilde{X}|Z]\}\ /\mathbb{V}[\tilde{X}]$, and

$$I_f(\tilde{X};\mathbb{E}[\tilde{X}|Z]) \geq \frac{1}{2}\log\left\{\mathbb{E}\left\{\mathbb{V}\left[\tilde{X}|Z\right]\right\}/\mathbb{V}\left\{\tilde{X}\right\}\right\} = \log\{\nu_f(\tilde{X}|Z)^{-1/2}\},$$

since $\mathbb{V}[\tilde{X}] = \mathbb{V}\{\mathbb{E}[\tilde{X}|Z]\} + \mathbb{E}\{\mathbb{V}[\tilde{X}|Z]\}$ by the law of total variance. This establishes the first inequality in Eq. 9 and, importantly, does so in a way that enables us to establish the second. It is a direct consequence of [16] that $\mathbb{V}\{\mathbb{E}[\tilde{X}|Z]\}/\mathbb{V}[\tilde{X}]$ is equal to $\sup_t \mathrm{Corr}^2(t(Z), \tilde{X})$, where the supremum is over all Borel measurable functions t such that $\mathrm{Corr}(t(Z), \tilde{X})$ is well defined. It follows that

$$1 - \nu_f(\tilde{X}|Z) = \mathbb{V}\{\mathbb{E}[\tilde{X}|Z]\}/\mathbb{V}[\tilde{X}] \geq \mathrm{Corr}^2(t(Z), \tilde{X}) \geq \mathrm{Corr}^2(Z, \tilde{X}), \tag{10}$$

for all $t(\cdot)$, which implies the second inequality in Eq. 9. ■

An analogue of Theorem 5 when the conditioning is on the response variable Z is given by the following Lemma. The proof is straightforward. A related lower bound is given without proof in the frequency domain by [24], the bound being on the mutual information rate in continuous-time.

Lemma 8. *Let (X, Z) be a random vector in $\mathbb{R}^d$, $d \geq 2$, with joint density with respect to Lebesgue measure, $f(x, z)$. Suppose there exists a one-to-one mapping $s : X \to \tilde{X}$ (whose domain is an open set and a support of X) such that $\tilde{X}$ has a Gaussian density. We assume that the mapping s is continuously differentiable with a Jacobian that is everywhere non-zero, and that $\nu_f(\tilde{X}|Z)^{-1/2}$ exists. Then, scaling by the reciprocal of the dimension of X,*

$$d_X^{-1} I_f(X;Z) \geq \frac{1}{2} d_X^{-1} \log\{\nu_f(\tilde{X}|Z)^{-1/2}\} \geq \frac{1}{2}\log\left\{\frac{[\det(\mathbb{V}[\tilde{X}])]^{d_X^{-1}}}{d_X^{-1}\mathrm{tr}(\mathbb{E}\{\mathbb{V}[\tilde{X}|Z]\})}\right\}_f, \tag{11}$$

and

$$C_{Z|X} \geq I_f(X;Z) \geq \log\left\{\nu_f(\tilde{X}|Z)^{-1/2}\right\}, \tag{12}$$

provided $f(x)$ is the density of an allowed input distribution, that is a distribution in $\mathcal{F}$.

Proof: A concise proof of the first inequality in Eq. 11 uses that $h(\tilde{X}|Z = z) \leq \frac{1}{2}\log[(2\pi e)^{d_X}\det(\mathbb{V}[\tilde{X}|Z = z])]$, which follows from the maximum entropy property of the multivariate Gaussian distribution for a given covariance matrix. Then $h(\tilde{X}|Z) \leq \frac{1}{2}\log[(2\pi e)^{d_X}\det(\mathbb{E}\{\mathbb{V}[\tilde{X}|Z = z])\}]$, where we have applied Jensen's inequality, using the concavity of the function $\log\{\det(\Sigma)\}$ for symmetric, non-negative definite matrices, Σ [25]. The result follows since $I_f(X;Z) = I_f(\tilde{X};Z) = h(\tilde{X}) - h(\tilde{X}|Z)$ and the entropy of the Gaussian $\tilde{X}$ is given by $h(\tilde{X}) = \frac{1}{2}\log[(2\pi e)^{d_X}\det(\mathbb{V}[\tilde{X}])]$. The second inequality in Eq. 11 follows because we take the covariance matrix $E\{\mathbb{V}[\tilde{X}|Z]\}$ to be positive definite and $0 < [\det(M)]^{1/m} \leq m^{-1}\mathrm{tr}(M)$ for a positive definite, $m \times m$ matrix, M. ■

The existence of the mapping $s(x)$ in Lemma 8 is not unduly restrictive. For example, when $d = 2$ and the input X has a strictly increasing distribution function taking values in $(0, 1)$, then $F_X(X)$ is uniformly distributed and can be invertibly transformed to a Gaussian random variable. Similar comments apply when $d > 2$, using the multivariate transformation of [26] to independent uniform r.v.'s on $(0, 1)$.

Rate distortion-type arguments using average mean-square error distortion and $\tilde{X}$ as the Gaussian source cannot establish Eq. 11 for $d_X > 1$. Such arguments [22, Section 4.5.2] show only that $\frac{1}{2}\log\{\det(\mathbb{V}[\tilde{\mathrm{X}}])^{d_X^{-1}}/d_X^{-1}\,\mathrm{tr}(\mathbb{E}\{\mathbb{V}[\tilde{X}|Z]\})\}_f$ is a lower bound for $d_X^{-1}I_f(X;Z)$ under the conditions of Lemma 8. As we have shown, our lower bound in Eq. 11, $\log\{\nu_f(\tilde{X}|Z)^{-1/2}\}$, is tighter than this one derived using average mean-square error rate distortion arguments with Gaussian source.

VI. Applications

The lower bounds on mutual information and capacity derived in previous sections will prove useful in at least two types of application: analysing the dependence between input and response vectors using empirical data; and analysing the information capacity of signaling and communication mechanisms for which physical models are available. An illustration of the first type of application using simulated data and further discussion are given immediately below. An existing example of the second type is given in [20] which examines the information capacity of optical fiber communication by employing a lower bound based on the Pearson correlation. Indeed, physical models of a communication mechanism can often be solved for their moments when distributional results are not feasible. For example, models of biomolecular signaling mechanisms are stochastic kinetic models of biochemical reaction networks [27] that can be solved approximately using system-size expansions of the master equation [28]. Such expansions can be used to provide fast, computational evaluation of $\mathbb{E}[Z|X]$ and $\mathbb{V}[Z|X]$ for many (rate) parameter vectors describing the network [29]. Theorem 5, Eq. 7 can then be used to approximate the capacity of the signaling mechanism and explore its parameter sensitivity.

A. Lower bound estimation

Estimation of mutual information rapidly becomes problematic as the dimension, d, of (X, Z) grows. A distinct advantage of the lower bounds in Eqs. 6 and 11 based on $\log(\nu^{-1/2})$ is that they are amenable to inference by using nonparametric regression to estimate the relevant conditional mean. Nonparametric regression and covariance matrix estimation methods for higher dimensions [30], [31] should break down more slowly than mutual information estimation methods as d grows. This is valuable for applications, including those in systems biology where multiple inputs and outputs often need to be considered. Estimation of our lower bounds should therefore prove useful for analysing the dependence between input and response in higher dimensions.

The following simulation study demonstrates that use of the lower bounds can substantially improve inference about mutual information when the sample size becomes limited for a given value of d. Here we use $d = 2$. Inferential procedures for the $d > 2$ setting lie beyond the scope of the present paper and will be explored in future work. The k-nearest neighbour point estimator [21], $\hat{I}_{knn}$, is widely regarded as the leading method for estimation of mutual information using continuously distributed data. We employ a *composite estimator* defined as the maximum of $\hat{I}_{knn}$ and the lower limit of our bootstrap confidence interval for $\log(\nu^{-1/2})$. This composite estimator makes use of our lower bounds to correct erroneous point estimates. We find that the lower bounds are able to provide substantial improvements to the downward bias and root mean square error (rmse) we report for the nearest-neighbour estimator.

We assume that we are given data $\{(X_i, Z_i); i = 1, ..., N\}$ for independent and identically distributed units and that the distribution of the input X is known (see the discussion following Eq. 1). We obtain confidence intervals, for example, for the lower bound in Eq. 9 based on $\nu(\tilde{X}|Z)$ as follows. (1) Obtain fitted values, $\hat{X}_i$, for the transformed, Gaussian input $\tilde{X}$ by nonparametric estimation of $\mathbb{E}[\tilde{X}|Z]$ using a smoothing spline; (2) Obtain the estimate $1 - \hat{\nu}_{\tilde{x}|z}$ as the ratio of the sample variance of $\hat{X}_i$ to the known variance of $\tilde{X}_i$ (see Eq. 10); (3) Obtain bias-corrected, accelerated ($\mathrm{BC_a}$) bootstrap confidence intervals [32] using the estimator $\log(\hat{\nu}_{\tilde{x}|z}^{-1/2})$. Details of the proposed procedure are given in the Appendix.

Figures 1 and 2 present simulation results for a range of true values of the mutual information and for two types of data generation mechanism: a bivariate normal distribution and a mixture model. In both, $Z = \alpha + \beta X + \varepsilon$, with ε normally distributed conditional on X (with constant variance independent of X). In the first, X has a marginal

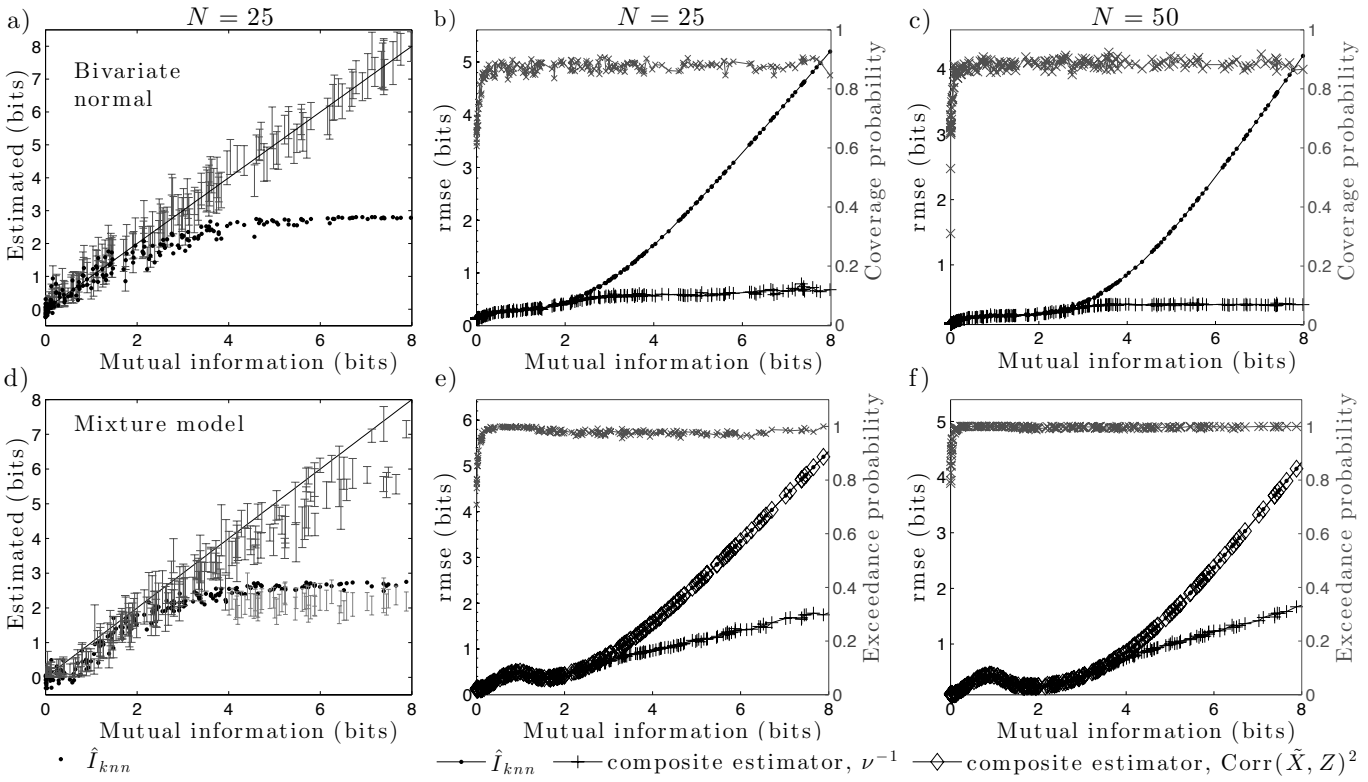

Fig. 1. **Bootstrap confidence intervals of the lower bounds $\log(\nu^{-1/2})$ can substantially improve inference about mutual information through use of a composite estimator.** Simulation results are shown for the bivariate normal (a–c), and for a mixture model incorporating the same linear regression model but with a mixture of normals distribution for X (d–f). N is the number of i.i.d. observations. a) and d): for $N = 25$, the $\mathrm{BC_a}$, nominally 90% confidence intervals for our lower bounds in Eqs. 6 and 9 resp., together with the k-nearest neighbour estimates, $\hat{I}_{knn}$, with $k = 3$. (Analogous confidence intervals for the lower bounds based on $\mathrm{Corr}(\tilde{X}, Z)^2$ are shown in d) for the mixture model when $I(X;Z) > 4$ bits). Parameter vectors for each model were sampled independently from their parameter spaces. A single data set is generated for each parameter vector in a) and d). Coverage probability (grey crosses, b) and c)) gives frequency with which the $\mathrm{BC_a}$ interval covers the true $I(X;Z)$; Exceedance probability (grey crosses, e) and f)) gives frequency with which $I(X;Z)$ exceeds the lower limit of the $\mathrm{BC_a}$ interval (nominally > 0.9). Root mean square errors (rmse) are plotted for $\hat{I}_{knn}$ (filled circles), and the composite estimators (see text) based on $\nu^{-1}_{Z|X}$ or $\nu^{-1}_{\tilde{X}|Z}$ (black crosses) and $\mathrm{Corr}(\tilde{X}, Z)^2$ (diamonds). Results based on 500 Monte Carlo replications.

normal distribution (under the data generating density, f), hence (X, Z) has a bivariate normal distribution, $X = \tilde{X}$, and the bounds $\log(\nu^{-1/2})$ in Eqs. 6 and 9 hold with equality. In the second, X is specified to be an equally-weighted mixture of 2 normals, and we obtain the pseudo-input $\tilde{X}$ by first transforming to uniformity using the probability integral transform and then transforming to normality. We adopt the second specification because $\mathbb{E}[\tilde{X}|Z]$ becomes non-linear (and sigmoidal), but the true value of $I(X;Z)$ is still known with precision through the use of a Monte Carlo average for $h(Z)$ (see Appendix). Details of the parameterisations of the models used are also given in the Appendix.

Panels a) and d) of Fig. 1 show our $\mathrm{BC_a}$ confidence intervals for the lower bounds based on Eq. 6 and 9 respectively, together with the point estimates $\hat{I}_{knn}$, for independently generated data sets corresponding to different true values of $I(X;Z)$ and for sample size $N = 25$. In [21], the authors recommend in practice to use values of k between 2 and 4. We therefore calculate $\hat{I}_{knn}$ using $k = 3$ nearest neighbours ($\hat{I}_{knn} = I^{(2)}(X, Z; k = 3)$ in [21]). The poor performance of $\hat{I}_{knn}$ with this sample size is evident for both models, particularly for mutual information in excess of 3 bits, where substantial, growing bias and rmse are evident (see Fig. 2 for plots of bias). Higher values of k result in worse bias and rmse of $\hat{I}_{knn}$ (not shown). The remaining panels of Fig. 1 depict, for sample sizes $N = 25$ and $N = 50$, various properties under repeated sampling: the frequency with which the $\mathrm{BC_a}$ interval for the lower bound covers [b) and c)] and has a lower limit exceeded by [e) and f)] the true mutual information; the rmse of $\hat{I}_{knn}$; and the rmse of our composite estimator, given by the maximum of $\hat{I}_{knn}$ and the lower limit of the $\mathrm{BC_a}$ interval. For both sample sizes, the nonparametric confidence intervals perform well under repeated sampling and provide substantial reductions in bias and rmse when comparing $\hat{I}_{knn}$ to the composite estimator (see also Fig. 2). Finally, in the mixture model where $\mathbb{E}[\tilde{X}|Z]$ is non-linear in Z, the lower bound based on $\mathrm{Corr}(\tilde{X}, Z)^2$ performs

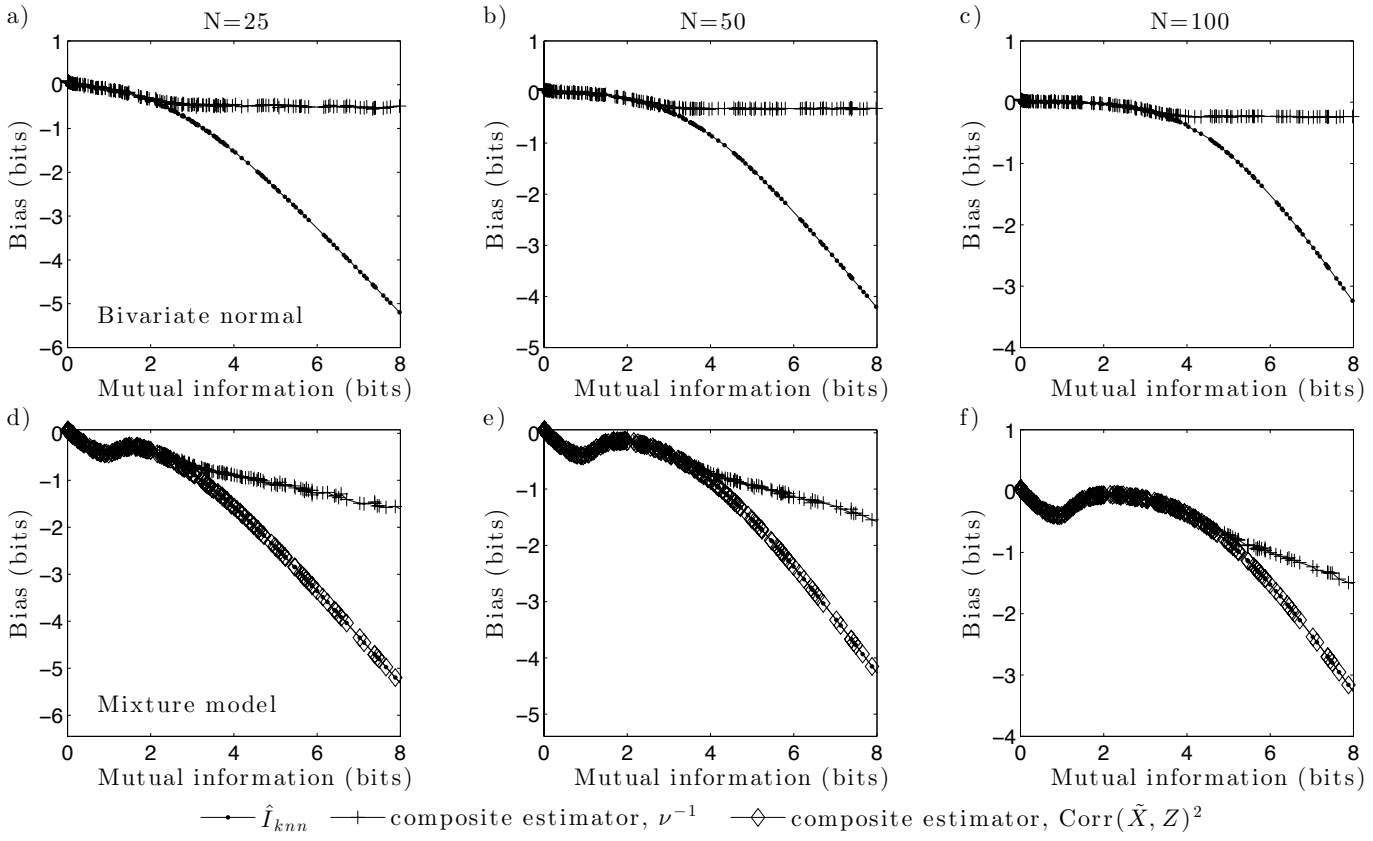

Fig. 2. **The lower bounds** $\log(\nu^{-1/2})$ **can substantially reduce bias through use of a composite estimator**. The lower bounds based on the Peason correlation do not reduce estimation bias for the mixture model with non-linear conditional mean. Biases are plotted for the k-nearest neighbour estimates, $\hat{I}_{knn}$, with $k = 3$ (filled circles), and for the composite estimators (see text) based on $\nu^{-1}_{Z|X}$ or $\nu^{-1}_{\tilde{X}|Z}$ (bivariate normal and mixture models respectively; crosses) and $\mathrm{Corr}(X, Z)^2$ (diamonds). Parameter vectors for each model were sampled independently from their parameter spaces. Results are shown for the bivariate normal (a–c), and for the mixture model (d–f). N is the number of *i.i.d.* observations in each data set. All results based on 500 Monte Carlo replications.

considerably worse than that based on $\nu_{\tilde{x}|z}$, as shown in panels d) to f) of Figs. 1 and 2. The corresponding $\mathrm{BC_a}$ intervals lie well below those based on $\nu_{\tilde{x}|z}$ and have lower limits below $\hat{I}_{knn}$ in all cases shown in panel d). The associated composite estimator consequently fails to reduce either the bias or the rmse of estimation.

VII. Appendix

A. Additional proofs

Proof: (Lemma 3) Let $g(x, z)$ be the multivariate Gaussian density with the same unconditional first and second moments as $f(x, z)$, and with marginal Gaussian density $g(x) = f(x)$. Thus, $\mathbb{V}_g[(X, Z)] = \mathbb{V}_f[(X, Z)]$. We use subscripts to identify the relevant joint density throughout. Notice that

$$I_f(X;Z) = \mathbb{E}_f\left\{\log\frac{g(X,Z)}{f(X)g(Z)}\right\} - \mathbb{E}_f\left\{\log\frac{g(X,Z)f(Z)}{f(X,Z)g(Z)}\right\} \geq \mathbb{E}_f\left\{\log\frac{g(X,Z)}{f(X)g(Z)}\right\},$$

where the second expectation of the equality is seen to be non-positive by applying Jensen's inequality and then integrating first with respect to x. Furthermore,

$$\mathbb{E}_f\left\{\log\frac{g(X,Z)}{f(X)g(Z)}\right\} = \mathbb{E}_g\left\{\log\frac{g(X,Z)}{f(X)g(Z)}\right\} = I_g(X;Z), \tag{13}$$

because $\mathbb{E}_f[\log\{g(\cdot)\}] = \mathbb{E}_g[\log\{g(\cdot)\}]$. For example, $\mathbb{E}_f[\log\{g(X,Z)\}] = \mathbb{E}_g[\log\{g(X,Z)\}]$ because

$$\mathbb{E}_f\left\{\begin{pmatrix} X-E[X] \\ Z-E[Z] \end{pmatrix}^{\mathrm{T}} \mathbb{V}_g[(X,Z)]^{-1} \begin{pmatrix} X-E[X] \\ Z-E[Z] \end{pmatrix}\right\} =$$

$$\mathrm{tr}\left\{\mathbb{V}_g[(X,Z)]^{-1}\mathbb{E}_f\left[\begin{pmatrix} X-E[X] \\ Z-E[Z] \end{pmatrix}\begin{pmatrix} X-E[X] \\ Z-E[Z] \end{pmatrix}^{\mathrm{T}}\right]\right\} = d,$$

since $\mathbb{V}_g[(X,Z)] = \mathbb{V}_f[(X,Z)]$. Evaluating $I_g(X;Z)$=$h_g(X) + h_g(Y) - h_g(X,Y)$ is straightforward since the marginal and joint densities under g are all Gaussian. We find

$$I_g(X;Z) = \frac{1}{2}\log\left\{\frac{\det(\mathbb{V}_g\,[\mathrm{X}])\det(\mathbb{V}_g\,[Z])}{\det(\mathbb{V}_g\,[(\mathrm{X},\mathrm{Z})])}\right\} = \frac{1}{2}\log\left\{\frac{\det(\mathbb{V}_f\,[\mathrm{X}])\det(\mathbb{V}_f\,[Z])}{\det(\mathbb{V}_f\,[(\mathrm{X},\mathrm{Z})])}\right\},$$

since g and f have identical second moments by construction. The stated results are then obtained by partitioning of the matrix $\mathbb{V}_f\,[(X,Z)]$. ∎

Proof: (Theorem 2). We have that $\nu(X_n|Z_n) = \nu(\breve{X}_n|Z_n) \to 0$. Since $\nu(\breve{X}_n|Z_n) = \mathbb{E}\{\mathbb{V}[\breve{X}_n|Z_n]\} = \mathbb{E}\{(\breve{X}_n - \mathbb{E}[\breve{X}_n|Z_n])^2\}$, it follows that $\breve{X}_n - \mathbb{E}[\breve{X}_n|Z_n]$ converges to zero in mean square (in L^2) and therefore $\breve{X}_n - \mathbb{E}[\breve{X}_n|Z_n] \to^{pr} 0$. Consider estimating $\breve{X}_n$ based on observation of Z_n as follows: the estimator $\hat{X}_n(Z_n)$ is equal to a point in the support of $\breve{X}_n$ which minimises the Euclidean distance from $\mathbb{E}[\breve{X}_n|Z_n]$. Let $x \in \mathbb{R}, \breve{x}_n \in \breve{\mathcal{X}}_n$ and notice that if $|\breve{x}_n - x| < \epsilon^*$, then $|\breve{x}_n - x| < \epsilon_n^* < |\breve{x}'_n - x|$, that is x is closer to $\breve{x}_n$ than to any other point $\breve{x}'_n$ in $\breve{\mathcal{X}}_n$. Therefore, if $|\breve{X}_n - E[\breve{X}_n|Z_n]| < \epsilon^*$, $\mathbb{E}[\breve{X}_n|Z_n]$ is closer to $\breve{X}_n$ than to any other point in the support, the estimator $\hat{X}_n(Z_n)$ is uniquely defined, and that estimator recovers $\breve{X}_n$ without error $\left(\hat{X}_n(Z_n) = \breve{X}_n\right)$. Thus, the probability of estimation error, p_{error}, satisfies

$$p_{\mathrm{error}} = P\left\{\hat{X}_n\,(Z_n) \neq \breve{X}_n\right\} \leq \mathbb{P}\left\{|\breve{X}_n - E[\breve{X}_n|Z_n]| \geq \epsilon^*\right\}.$$

Since $\breve{X}_n - \mathbb{E}[\breve{X}_n|Z_n] \to^{pr} 0$, p_{error} must therefore tend to zero as $n \to \infty$. Fano's Inequality gives

$$H(p_{\mathrm{error}}) + p_{\mathrm{error}}\log|\mathcal{X}_n| \geq H(\breve{X}_n|Z_n) = H(X_n|Z_n),$$

since $|\breve{\mathcal{X}}_n| = |\mathcal{X}_n|$ and the rescaling does not change the conditional entropy. Therefore $H(X_n|Z_n) \to 0$ as $n \to \infty$. ∎

Models, parametrisations and algorithms used in the simulation study of Section VI-A

Figures 1 and 2 present simulation results for two data generation mechanisms. In both, $Z = \beta X + \varepsilon$ with $\varepsilon \sim N(0, \sigma_\varepsilon^2)$ and ε independent of X. The two models, together with the schemes used to generate parameter vectors for the results shown in Figures 1 and 2, are as follows:

1) *Bivariate normal model*: $X \sim N(0, \sigma_x^2)$. Model parameters were sampled as follows: i) β uniformly distributed on $(1, 10)$; ii) $\sigma_\varepsilon^2 = 10^{\theta_1}$ with θ_1 uniformly distributed on $(-2, 2)$; and iii) $\sigma_x^2 = 10^{\theta_2}$ with θ_2 uniformly distributed on $(-2, 2)$.
2) *Mixture model*: X is an equally-weighted mixture of 2 normal distributions, that is $f_x(x) = \frac{1}{2}N(\mu_1, \sigma_1^2) + \frac{1}{2}N(\mu_2, \sigma_2^2)$. Model parameters were sampled as follows: i) $\beta = 10^{\theta_1}$, with θ_1 uniformly distributed on $(-1, 1)$; ii) $\sigma_\varepsilon^2 = 10^{\theta_2}$ with θ_2 uniformly distributed on $(-2.5, 2.5)$. We set $\mu_1 = -\mu_2 = 5$ and $\sigma_1^2 = \sigma_2^2 = 25/4$.

The mixture model allows precise evaluation of $I(X;Z) = h(Z) - h(Z|X)$ via Monte Carlo sampling. We have $h(Z|X) = \frac{1}{2}\log(2\pi\mathrm{e}\sigma_\varepsilon^2)$. Note that the marginal density $f(Z)$ is also an equally-weighted mixture of 2 normals which we can express in closed form. Hence, we can also estimate $h(Z) = -\mathbb{E}[\log f(Z)]$ as the Monte Carlo average of $\log f(Z_m)$ where Z_m $(m = 1, ..., M)$ is a draw from the mixture model. For our numerical calculations we set $M = 10^5$, and monitored convergence of the Monte Carlo average.

Computations were implemented in R (version 2.12.2). The non-parametric estimation of $\mathbb{E}[\tilde{X}|Z]$ was performed using the 'smooth.spline' function (an implementation of smoothing splines [33]) with the number of knots set to 10; the smoothing parameter was chosen using cross-validation on the original dataset; all other parameters were set to their default values. 90% $\mathrm{BC_a}$ confidence intervals were calculated from $B = 2000$ bootstrap replications (using the 'boot' package). For the k-nearest neighbour estimation of mutual information [21] we used the authors' 'MIxnyn' function within their MILCA suite (available at http://www.klab.caltech.edu/~kraskov/MILCA/).

REFERENCES

[1] Linfoot, E. H. (1957) *Information and Control* **1**, 85–89.
[2] Joe, H. (1989) *Journal of the American Statistical Association* **84**, 157.
[3] Brillinger, D. and Guha, A. (2007) *Journal of Statistical Planning and Inference* **137**, 1076–1084.
[4] Reshef, D. N., Reshef, Y. A., Finucane, H. K., and Grossman, S. R. et al. (2011) *Science* **334**, 1518–1524.
[5] Shannon, C. E. (1948) *Bell System Technical Journal* **27**, 379–423.
[6] Brennan, M. D., Cheong, R., and Levchenko, A. (2012) *Science* **338**, 334–5.
[7] Rieke, F., Warland, D., de Ruyter vanSteveninck, R., and Bialek, W. (1999) Spikes: Exploring the Neural Code, MIT Press.
[8] Cheong, R., Rhee, A., Wang, C. J., Nemenman, I., and Levchenko, A. (2011) *Science* **334**, 354–358.
[9] Bowsher, C. G. and Swain, P. S. (2012) *Proceedings of the National Academy of Sciences USA* **109**, E1320–8.
[10] Bowsher, C. G., Voliotis, M., and Swain, P. S. (2013) *PLoS Computational Biology* **9**, e1002965.
[11] Panzeri, S., Senatore, R., Montemurro, M. A., and Rasmus, S. et al. (2012) *Journal of Neurophysiology* **98**, 1064–1072.
[12] Prelov, V. and Verdú, S. (2004) *IEEE Transactions on Information Theory* **50**, 1567–1580.
[13] Guo, D., Shamai, S., and Verdú, S. (2005) *IEEE Transactions on Information Theory* **51**, 1261–1282.
[14] Guo, D., Shamai, S., and Verdú, S. (2008) *IEEE Transactions on Information Theory* **54**, 1837–1849.
[15] Cover, T. M. and Thomas, J. A. (2006) Elements of Information Theory, John Wiley & Sons Inc., second edition.
[16] Rényi, A. (1959) *Acta Mathematica Academiae Scientiarum Hungarica* **10**, 441–451.
[17] Bernardo, J. M. (1979) *The Annals of Statistics* **7**, 686–690.
[18] Bernardo, J. M. and Smith, A. F. M. (2000) Bayesian Theory, John Wiley & Sons Inc.
[19] Dawid, A. P. (2010) Seeing and doing: The Pearlian synthesis. In: Dechter, R., Geffner H., and Halpern J. Y. eds. *Heuristics, probability and causality: A tribute to Judea Pearl*: College Publications, 309–325.
[20] Mitra, P. P. and Stark, J. B. (2001) *Nature* **411**, 1027–30.
[21] Kraskov, A., Stögbauer, H., and Grassberger, P. (2004) *Phys Rev E* **69**, 1–16.
[22] Berger, T. (1971) Rate-Distortion Theory, Prentice Hall.
[23] Gebelein, H. (1941) *Zeitschrift für angew* **21**, 364–379.
[24] Rieke, F., D., W., and Bialek, W. (1993) *Europhysics Letters* **22**, 151–156.
[25] Cover, T. M. and Thomas, J. A. (1988) *SIAM Journal on Matrix Analysis and Applications* **9**, 384–392.
[26] Rosenblatt, M. (1952) *The Annals of Mathematical Statistics* **23**, 470–472.
[27] Bowsher, C. G. (2010) *The Annals of Statistics* **38**, 2242–2281.
[28] Grima, R. (2010) *The Journal of Chemical Physics* **133**, 035101–035101.
[29] Thomas, P., Matuschek, H., and Grima, R. (2012) *PloS One* **7**, e38518.
[30] Friedman, J. H. (1991) *The Annals of Statistics* **19**, pp. 1–67.
[31] Bickel, P. J. and Levina, E. (2008) *The Annals of Statistics* **36**, pp. 199–227.
[32] Efron, B. and Tibshirani, R. J. (1994) An Introduction to the Bootstrap, Chapman & Hall/CRC.
[33] Green, P. J. and Silverman, B. W. (1994) Nonparametric Regression and Generalized Linear Models, Chapman & Hall/CRC.